

STATISTICS WORKSHOP II

United States Department of Agriculture

*Exploratory
and
Confirmation
Data
Analysis*

presented by
Vicki Lancaster
NEPTUNE AND CO., INC.
vlancast@neptuneandco.com

Data Analysis Introduction

What is data,

A collection of numerical values recording the magnitudes of various attributes of the objects under study.

(Hand, 1999)

What is data analysis,

The processing of the data.

(Hand, 1999)

Why do we need data?

Data Analysis *Introduction*

In God we trust,



all others must bring **data**.

Unknown

Data Analysis Introduction

Data analysis is *not* a case of simply applying a directory of tools to a given problem.

critical assessment,
exploration,
testing, and
evaluation.

It is a domain that requires *intelligence* and *knowledge* and *expertise* about the data. It is a challenging and demanding discipline.

It is a discipline that which is continuing to evolve.

(Hand, 1999)

Data Analysis Introduction

There are two broad categories of data analysis: *exploratory* and *confirmatory*.

- 1. *Exploratory data analysis* (EDA) is concerned with **searching** for clues and finding evidence.
- 2.2. *Confirmatory data analysis* (CDA) (CDA) is concerned with **evaluating** the evidence.

SESSION OUTLINE	
<i>Data Analysis</i>	
EDA	CDA
<i>Four Themes of EDA</i>	<i>Goodness-of-Fit Tests</i>
1. <i>resistance</i>	1. <i>chi-square</i>
2. <i>residuals</i>	2. <i>EDF</i>
3. <i>re-expression</i>	3. <i>moment</i>
4. <i>displays</i>	4. <i>regression</i>

EDA *Introduction*

What is *exploratory data analysis* (EDA)?

EDAEDA is a process that uses non-EDA is a process that suchsuch as graphical methods, to gain insight into asu data.

*I t It character It characterizes It characterizes It characterizes
and modeling are driven by data.*

If you think of your data set as a story written in numbers, then EDA is the story written in pictures.

EDA methods are used:

to *isolate* patterns and relationships,

to *uncover* unexpected behavior,

to *confirm* or *disprove* or assumptions, and

to *reveal* information.

EDA Introduction

Why is *exploratory data analysis* important?

Most classical procedures are based on *assumptions* about the characteristics of a variable, and of a variable, the analyses depends upon the validity of the *assumptions*.

The *graphical methods* of EDA provide powerful diagnostic tools for confirming assumptions are not met, for suggesting corrective actions.

EDA Introduction

For example,

if you if you bif you blindly conducted a *one sample t-test* that looked like ...

you would fail to reject the null hypothesis.

EDA Introduction

But,
if you had done a few EDA plots on the data first ...

you would have noticed a potential outlier.

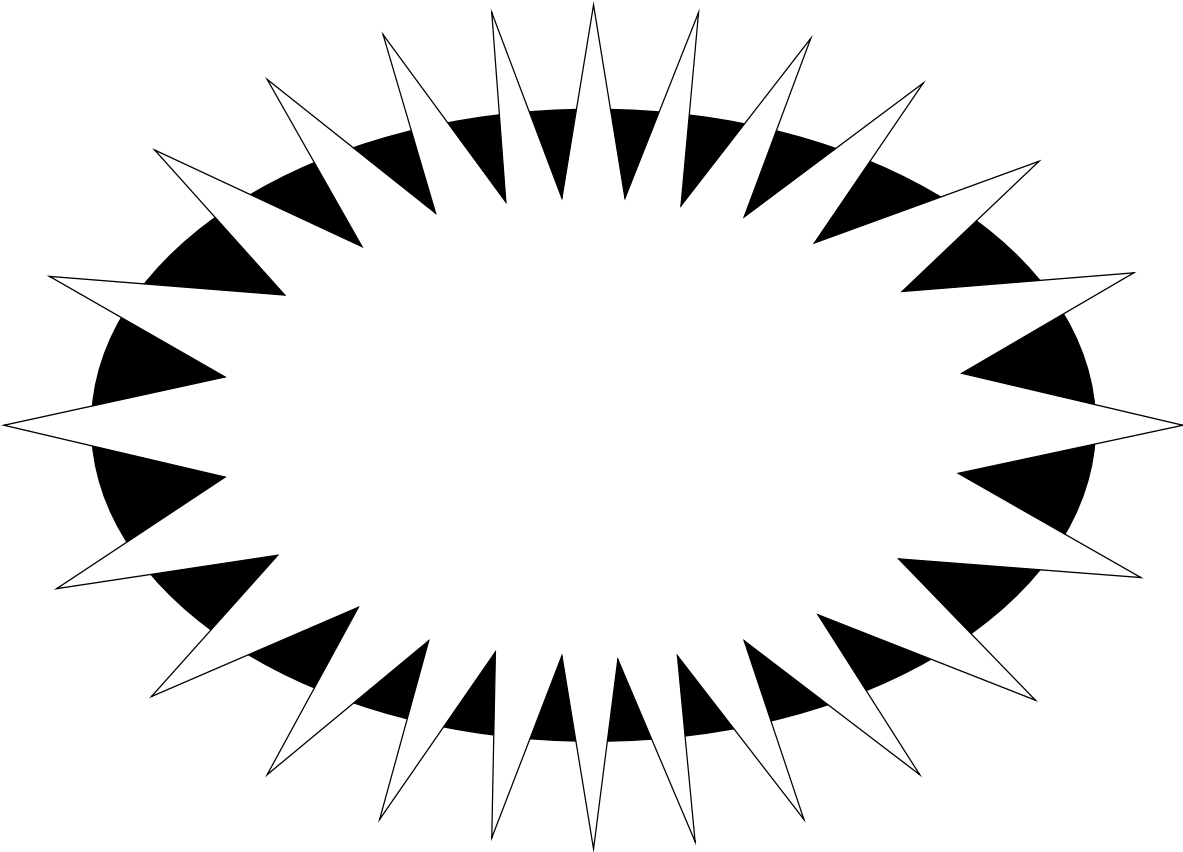
EDA Introduction

If,

it turns out the outlier is due to
example a data entry error) and it is corrected ...

EDA Introduction

The results of your t-test are now different.



Graphical Methods for Data Analysis

John M. Chambers

William S. Cleveland

Beat Kleiner

Paul A. Tukey

EDA Introduction

The first published presentation of EDA (1970 - 1977) was the preliminary Tukey. His 1977 book,

Tukey, J. W. (1977). *Exploratory Data Analysis*. Addison-Wesley, Reading Mass.

represents the definitive account on the subject. It wanted to dispel the *myth* that we are not allowed to look at the data prior to modeling.

At that time there was a tension between competing points of view:

- that a hypothesis must not be data driven; and

- that EDA was needed prior to inferential statistics to understand how rich the data are to support.

Four Themes of EDA

Resistance

Residuals

Re-expression

Displays

Four Themes of EDA

Resistance is a term used to denote a property of measures of location or spread that is relatively unaffected by the presence of outliers. The median is an example of a resistant measure of location and the interquartile range (IQR) is an example of a resistant measure of spread.

A summary statistic is *resistant* if

it is insensitive to any small change of the data, and

to any large change in a small part of the data.

Four Themes of EDA

Robust versus Resistance

Robust is used to describe an *inference procedure* that is stable when model assumptions are violated. For example, the *t-test* is *robust* with respect to the assumption of normality.

Robustness sensitivity to model assumptions.

Resistance is used to describe a *statistic* that is arithmetically stable under data values. For example, the *median* is a *resistant* estimator.

Resistance sensitivity to the data.

Four Themes of EDA

Resistant summary statistics:

pay attention to the *main body* of the data;
given little attention to outliers; and

are useful in graphical methods;
the construction of box plots.

For example, look at the
9 and 10. Are any of these *resistant*?

Four Themes of EDA

EDAEDA charactEDA characteriEDA characterized
decomposing the data into structure and noise,

$$data = \text{fit} + \text{residuals},$$

[illegible]

This process has its roots in the paradigm of partitioning variability into parts, explained and unexplained, and the notion simply uses that only on the observed treatment possible.

Four Themes of EDA

The philosophy of EDA is that data is not complete without a careful examination of the *residuals*.

Resistant analyses provide dominant behavior and unusual behavior in data.

Residuals contain any drastic departures from the expected pattern, as well as random fluctuations.

Four Themes of EDA

Re-expression involves the question of what scale would help to simplify the analysis of the data.

Re-expression into another scale may help to

achieve symmetry,
facilitate interpretation,
promote constancy of variance,
achieve a more linear relationship, or
simplify structure for two-way tables.

depending on the structure of the data.

Re-expression most often comes from the family of functions known as power functions, which take y into y^n , together with the logarithm.

Four Themes of EDA

Ladder of Transformations: $y = x^p$

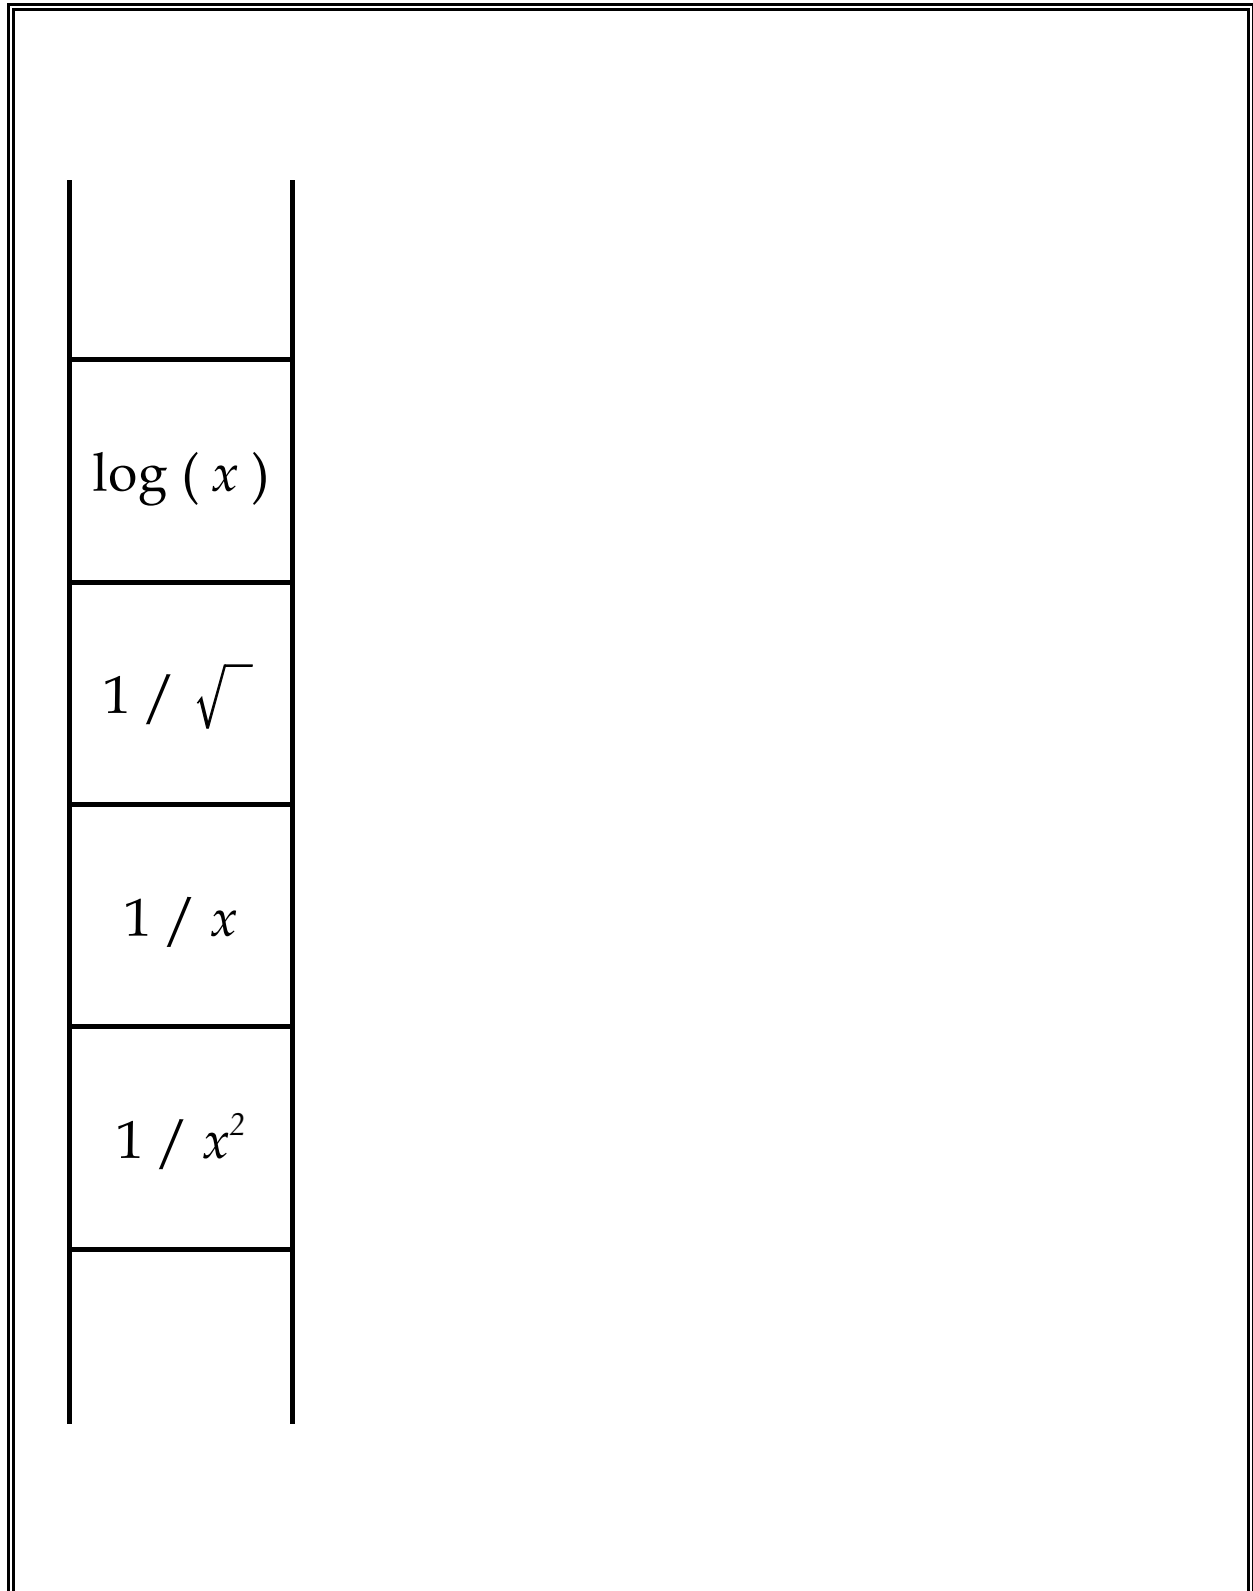
$p = 2$	x^2	square
$p = 1$	x	(none)
$p = 1/2$	$\sqrt{}$	square root
$p = 0$	$\log (x)$	log
$p = 1/2$	$1 / \sqrt{}$	reciprocal square root
$p = 1$	$1 / x$	reciprocal
$p = 2$	$1 / x^2$	reciprocal square

Four Themes of EDA

In conjunction with the ladder of transformations suggestions for transforming to achieve normality.

x^2
x
$\sqrt{}$
$\log (x)$
$1 / \sqrt{}$
$1 / x$
$1 / x^2$

Four Themes of EDA



Four Themes of EDA

Displays meet the need to see the meet the need to see the behavior to reveal the unexpected features, such as *outliers*; and confirm or disprove *assumptions*, such as the distributional assumptions of normality.

Tukey argued that (Tukey argued that (go Tukey argued that allow unexpected values allow unexpected values to pre-identified, identified, identified, identified, models can be explained for them.

Displays of EDA

What are *outliers*?

An *outlier* is defined as any value of the variable that falls outside the pattern of the other values. Exactly what is subjective. There is no quantitative rule that can be generalized to all identifying outliers.

Displays of EDA

Why is the detection of *outliers* important?

Outliers can greatly influence the value of the mean statistic and the conclusion of the test.

In many instances there is an associated outlier(s) such as an input error.

But even when there is no assignment, the data value is not changed or omitted.

Kruskal (1960a) *doctrine* states:

it is of great importance to preach the it is of great importance to pr
outliers should always be reported, even be reported, even when one
their causes are their causes are known or when one rejects them for their ca
good rule or reason. The immediate pregood rule or reason. The im
statistical analysis are alms statistical analysis are almos statistical a
suppressing announcement of observations that suppressing anno
pattern; we must maintain a st pattern; we must maintain a str pat
pressures.

Displays of EDA

There are two principal methods for dealing with outliers: *identification* and *accommodation*. If an outlier(s) is detected or identified by one of several ways (Tietjen, 1986):

- omit the outlier(s) and treat as a new sample;

- omit the outlier(s) and censor the sample;

- ask the experimenter to replace (verify) the outlier;

- Winsorize the outliers, i.e., replace the value of the nearest good observation;

- present all analyses with and without outlier(s). When there are no differences between the two sets of analyses, both results should be the same. When there is a difference, both results should be presented for conclusions.

Displays of EDA

What are *assumptions*?

Assumptions are the rules under which the rules and conclusions drawn from applying an inferential method are valid.

The implicit *assumptions* of an inferential method are the rules that govern its use.

These methods are trustworthy if the rules that govern their use are met.

Displays of EDA

Why is it important to evaluate *assumptions*?

The distributional assumptions of the data will in turn determine what inferential statistical method is appropriate, parametric or non-parametric.

Estimation procedures such as the calculation of confidence intervals, to prediction intervals depend strongly on the underlying distribution. When a distribution is assumed extreme tail percentiles are often lognormal. This is important in environmental work where data are often lognormal.

Displays of EDA

These *assumptions* can be viewed as fallacious categories:

- 1.1. those that constrain *how* the data are collected as the requirement of a random sampling process and the use of an appropriate random process for assigning treatment (determined by the experimental design); and
2. those that constrain the *characteristics* of the data such as a requirement that the data be normally distributed.

An experimental design is adopted *prior* to data collection to assure the results generated by a random (or a restricted random) process.

EDA techniques are employed to evaluate the characteristics of the data.

Displays of EDA

A major contribution of the development of EDA has been the emergence of a wide variety of new graphical techniques.

Commonly used plots to check for distributional assumptions and outliers include,

histograms,

box-plots, and

quantile - quantile plots (*Q-Q plots*).

Displays of EDA

Definitions . . .

A *boxplot* is a rectangle, the top and bottom of the rectangle represent the upper and lower quartiles of the data, the horizontal line inside the rectangle represents the median. The vertical lines (whiskers) extend from the box to the minimum and maximum values not beyond a *standard span* from the quartiles. These lines are often drawn individually.

The *standard span* is 1.5 Inter-Quartile Range (IQR).

Displays of EDA

Definitions . . .

The *quantile* of the data is a value that divides the data into two groups, so that a certain percentage of observations fall below the quantile. For example, the 75th quantile ($Q(.75)$) divides the data set such that three fourths of the observations fall below and one fourth fall above.

Note: The *median* is the 50th quantile, $Q(.50)$.

The *upper quartile* is the 75th quantile, $Q(.75)$.

The *lower quartile* is the 25th quantile, $Q(.25)$.

The $IQR = Q(.75) - Q(.25)$.

Displays of EDA

Figure 1. Data from a
Normal Distribution, Sample Size=38

Displays of EDA

Figure 2. Data from a
Lognormal Distribution, Sample Size=38

Displays of EDA

The box plot is a visual display of the *five-number summary* of a data set.

Definition . . .

The *five-number summary* of a data set consists of the smallest observation (lower whisker), the lower quartile (bottom line in the box), the upper quartile (top line in the box), and the largest observation (upper whisker), written in order from smallest to largest.

For example, the *five-number summary* for the data displayed in Figure 1 is:

<i>Med.</i>	
<i>Q(.25)</i>	<i>Q(.75)</i>
<i>Min.</i>	<i>Max.</i>

4.30	
3.30	6.45
0.51	8.09

Diplays of EDA

It is interesting to compare the descriptive statistics of the data displayed in Figures 1 and 2.

<i>Descriptive Statistics</i>	<i>Fig. 1 Normal</i>	<i>Fig. 2 Lognormal</i>
Lower Quartile	3.30	33.72
Mean	4.68	768.50
Median	4.30	168.70
Upper Quartile	6.45	573.20
Standard Deviation	2.06	1828.52
Range [Min. Max.]	7.58	10116.74
IQR [Q(.75) Q(.25)]	3.15	539.48
Coefficient of Skewness	0.14	4.08
Coefficient of Kurtosis	2.12	20.13

Displays of EDA

Definitions . . .

The third moment about the mean is a measure of asymmetry called skewness. Symmetric distributions will have a skewness of 0. Distributions that are skewed to the left have a skewness < 0 , and distributions that are skewed to the right will have a skewness > 0 .

Of interest is the *standardized third moment* or the *coefficient of skewness*, $\sqrt[3]{\frac{m_3}{m_2^3}}$,

where,

$$\sqrt[3]{\frac{m_3}{m_2^3}}, \text{ and}$$
$$\frac{m_3}{m_2^3}.$$

Displays of EDA

Definitions . . .

The fourth moment about the mean of curvature or *kurtosis*, which is the measure of flatness of a density near its center.

Of interest is the *standardized fourth moment coefficient of kurtosis*, b_2 ,

where,

,

—

, and

—

.

Displays of EDA

Values of $\sqrt{b_1}$ and b_2 close to 0 and $3(n-1)/(n+1)$, respectively indicate normality.

Values differing from these are indicators of non-normality.

The signs and magnitude of this information on the type of non-normality.

$\sqrt{b_1} > 0$ positively (or right) skewed,

$\sqrt{b_1} < 0$ negatively (or left) skewed,

$b_2 > 3(n-1)/(n+1)$ relates to heavier tails than the normal, and

$b_2 < 3(n-1)/(n+1)$ relates to lighter tails than the normal.

Displays of EDA

Some observations on the Some observations on the desc
Normal distribution:

mean median,

skewness 0,

kurtosis 3.

Some observations on the descriptive Some observations
Lognormal distribution:

mean >> median,

skewness >> 0,

kurtosis >> 3.

Displays of EDA

Definitions . . .

A *histogram* partitions the range of data into several nonoverlapping intervals, called bins, and counts the number of observations in each bin. The number of counts in each bin can be displayed on a density scale, where the height of the bar represents the probability density, or on a frequency scale, where the height represents the bin counts. The histogram is completely determined by two parameters, the *bin width* and the *bin origin*.

Note: The *histogram* is the simplest and most familiar example of a probability distribution function).

Displays of EDA

Figure 3. Normal and Lognormal Data
from Figures 1 and 2 Displayed Using Histograms

Displays of EDA

What can we say about histograms in Figure 3?

Figure (a) appears to be bimodal.

Figure (b) is definitely right skewed.

Displays of EDA

Note: Histograms can give different visual impressions that are a function of the arbitrary choice of the intervals. The choice determines whether we retain smoothness and simplicity (a) or show more detail (b).

Displays of EDA

The tradeoff between smoothness and closeness to the data is determined by the thumb rule for determining the bin width (h) where the reference density is the Normal are:

$$h_1 = \{\text{range}(\mathbf{y}) / \log_2 n + 1\}, \text{ Sturges formula,}$$

$$h_2 = \{3.5 n^{-1/3}\}, \text{ Scott (1979), and}$$

$$h_3 = \{2 \text{ IQR } n^{-1/3}\}, \text{ Freedman \& Diaconis}, \text{ Freedman \&}$$

where \mathbf{y} is the sample vector and s is the estimated standard deviation of \mathbf{y} .

Note: Take a look at the web page
<http://www.stat.sc.edu/~west/javahtml/Histogram.html>
for an applet on histograms and bin width.

Displays of EDA

Herbert Sturges (1926) was the first to propose systematic guidelines for designing a histogram. He observed that the binomial distribution, $B(n, p)$, could be used as a model of an optimally constructed histogram.

Construct a frequency histogram with width 1 centered on the point $i = 0, 1, 2, \dots, k - 1$.

Choose the bin count c . Choose the binomial coefficient

.

The total sample size is

.

By the Binomial expansion Sturges rule follows

$$k = 1 + \log_2 n.$$

Displays of EDA

In the case of Scott and F&D, the rule compromises between histogram using the Normal as the reference distribution. The bin width, h , is viewed as a smoothing parameter.

The *variance* can be reduced by making the bins are wide and approximate height. The *variance* can be eliminated by choosing $h = \text{range}(\mathbf{y})$ (all observations are 1).

The *bias* can be reduced by making bins are narrow. The bias can be eliminated by choosing $h = \{\min |y_i - y_j|, \text{wh}\}$, where $i \neq j$ (each observation is its own bin).

Displays of EDA

Definition . . .

If the mean of all possible values of a statistic is equal to a parameter, the statistic is called an unbiased estimator of that parameter.

For example,

The sample mean, \bar{x} , is an unbiased estimator of the population mean, μ , because the means of all possible samples of a given size are equal to the population mean.

Displays of EDA

Definition (cont.) . . .

Example of the Mean as an Unbiased Estimator

Population of 3 observations (2, 4, 6) where $\bar{x} = 4$

Sample No.	All possible samples (x_1, x_2)	Mean $\bar{x} = (x_1 + x_2)/2$
1	2, 2	2
2	2, 4	3
3	2, 6	4
4	4, 2	3
5	4, 4	4
6	4, 6	5
7	6, 2	4
8	6, 4	5
9	6, 6	6
Sum		36
Mean	$\bar{x} / 9$	4

Displays of EDA

Definition . . .

Let $y_1, y_2, y_3, \dots, y_n$ be a random sample from a probability distribution depends on an unknown parameter, θ . Let $T = f(y_1, y_2, y_3, \dots, y_n)$ be a statistic (for example, statistic (squared error (MSE) of T , as an estimator for θ is

$$\text{MSE}(T) = \text{Var}(T) + [\text{Bias}(T)]^2,$$

where $[\text{Bias}(T)]^2 = [E(T) - \theta]^2$ and E represents the *expected value*.

Displays of EDA

The *bias* and *variance* are controlled by choosing an intermediate value between $\{\min |y_i - y_j|, \text{range}(\mathbf{y})\}$ and allowing the bin width w to decrease as the sample size increases.

Definition . . .

A density estimator is said to be an *optimal density estimator* if the mean square error of the estimator tends to zero as $n \rightarrow \infty$.

An *optimal smoothing parameter*, h^* , is defined to be that choice that minimizes the MSE.

Displays of EDA

Table 1. Comparison of the Number of Bins from the Three Normal Reference Rules*

<i>Number of Bins</i>	<i>Sturges</i> $[\log_2 n + 1]$	<i>Scott</i> $\left[\frac{3.7}{n^{1/3}} \right]$	<i>F&D</i> $\left[\frac{10}{\sqrt{n}} \right]$
50	5.6	6.3	8.5
100	7.6	8.0	10.8
500	10.0	13.6	18.3
1,000	11.0	17.2	23.2
5,000	13.3	29.1	39.6
10,000	14.3	37.0	49.9
100,000	17.6	79.8	107.6

** Scott, D. W. (1992) *Multivariate Density Estimation. Multivariate Density Estimation and Visualization*. New York: John Wiley & Sons. The data used to estimate the bin numbers were generated by a multivariate normal distribution with mean vector (0, 0, 0) and covariance matrix (3).

Displays of EDA

Some observations about Table 1:

The rules are comparable for sample sizes greater than 100.

For sample sizes greater than 100, the rules will provide an oversmoothed histogram and waste much of the information in the data.

The Freeman-Diaconis rule has a larger bin width than Scott's rule and therefore produces a smoother histogram.

Displays of EDA

Comparison of the Number of
Bins from the Three Normal Reference Rules
for the Data in Figures 1 and 2

<i>Number of Bins</i>	<i>Sturges</i>	<i>Scott</i>	<i>F&D</i>
Normal	7	4	5
Lognormal	7	6	32

Displays of EDA

A *cool* example that illustrates the power of EDA.

For the LANL Environmental Restoration For the LANL Environmental
extensive site characterization was performed.

Surface soil samples are collected and compared to
LANL background data. LANL background data
samples were collected from MDAs. Samples were collected
elemental uranium.

Displays of EDA

MDA G Uranium Concentration (mg/kg)

Displays of EDA

What can we see in these data? What can we

Contaminations - different historical disposal operations.

Geology issues - different background concentration

Chemistry issues - different analytical methods, data comparability issues.

Sample collection issues - different sampling methods, different field team, issues.

Displays of EDA

MDA G Uranium Concentration (mg/kg)
by Analytic Technique

Displays of EDA

The problem was one of lack of inclusion of two very different analytical methods.

EDA plots brought about a change in policy, KPA EDA no longer being used LANL ER Project.

Displays of EDA

Definitions . . .

AA theoretical quantile-quantile plot (Q-Q plot) or probabilityprobability plot is ob is obtained is quantilesquantiles of the observed data against the correspondingcorresponding quantiles ofcorrespon distribution (for example, the normal).

Displays of EDA

How do you construct a normal Q - Q plot?

Let $y_1, y_2, y_3, \dots, y_n$ represent the raw data:

sort the data from smallest to largest
 $y_{(3)}, \dots, y_{(n)}$,

calculate the empirical quantiles
observation, $Q_e(p_i)$, where

$$p_i = (i - 0.5)/n$$

((i.e., for a sample size of $n = 20$, the fifth
observation, $y_{(5)}$, is the $\{(5 - 0.5)/20\}^{\text{th}}$ quantile
 $Q_e(0.225)$),

calculate the corresponding quantiles for the
standard normal distribution ($z = 0$, $z = 1$), if $z = 1$, i.e.,
cumulative distribution function of the standard
normal, then

$$Q_t(p_i) = F^{-1}(p_i).$$

Displays of EDA

Let's try an example.

Table 2. Simple example for constructing a Standard Normal Q - Q plot

$y_{(i)}$	$p_{(i)}$	$Q_e(p_i)$	$Q_t(p_i)^*$
7	0.05	7	1.64
8	0.15	8	1.04
11	0.25	11	0.38
13	0.35	13	0.13
14	0.45	14	0.13
17	0.55	17	0.38
18	0.65	18	0.38
19	0.80	19	0.84
19	0.80	19	0.84
20	0.95	20	1.64

* These values are found by looking them up in a table or using a software package that calculates quantiles.

Displays of EDA

If the quantiles of the empirical distribution and the quantiles of the theoretical distribution are plotted on a straight line then the distributions are similar.

We can think of this another way.

Let,

$$F(y) = \left(\frac{y - \mu}{\sigma} \right) = G(z)$$

where,

$z = \frac{y - \mu}{\sigma}$ is the standardized variable and

$G(z)$ is the CDF of the standard normal variable Z .

$$z = G^{-1}(F(y)) = \frac{y - \mu}{\sigma} = \frac{y - \bar{y}}{s}$$

or in terms of y on z ,

$$y = \bar{y} + z \cdot s$$

What we are doing is transforming the values to standard normal variates. In a Q plot the intercept is \bar{y} and the slope is s .

Displays of EDA

Constructing the Standard
Normal Q-Q plot from the data in Table 2.

Displays of EDA

The Standard Normal Q-Q plot for the data

Displays of EDA

Properties of the theoretical Q - Q plot:

If the theoretical distribution is a good approximation to the empirical distribution, the points on the plot will fall near the $y = x$ line.

If the points follow a line that is not the $y = x$ line, then the appropriate positive or negative constant could be added to all data points to shift the configuration onto the $y = x$ line.



Conclude: The empirical distribution is compatible with the theoretical distribution if they have different means or medians.

Displays of EDA

Properties of the theoretical Q - Q plot (cont.):

If the points follow a line that is nearly straight and pass through the origin and the $y = x$ line, then it is possible to find an appropriate positive constant by which to multiply all observations to multiply all observations vertically and shift the configuration onto the $y = x$ line.



Conclude: The empirical distribution is compatible with the theoretical distribution if they have different spreads (standard deviation or interquartile range).

Displays of EDA

Properties of the theoretical $Q-Q$ plot (cont.):

The straightness of the theoretical $Q-Q$ plot is a good judge whether the empirical and theoretical distributions have the same distributional shape. Shifts and tilts away from the straight line indicate differences in location and spread, respectively.

A single theoretical $Q-Q$ plot compares a set of data not just to one theoretical distribution but simultaneously to a whole family of distributions with different locations (means) and spreads (standard deviations).

Displays of EDA

How do we interpret a $Q-Q$ plot?
are deviations from the straight line pattern?

Not only does the $Q-Q$ plot provide a warning provide
match is poor, but it may also mismatch is poor, but it may al
mismatch.

When there are departures from linearity in a $Q-Q$ plot
they frequently match one of the following
descriptions:

1. outliers at either end,
2. curvature at both ends,
3. convex or concave curvatures, and
4. horizontal segments, plateaus, or gaps.

Displays of EDA

Departures from linearity:

1. *outliers at either end,*

Are the most extreme observations even larger than could be reasonably expected for sample size from the distribution in question?

The theoretical *Q-Q plot* provides an informative answer.

Displays of EDA

Departures from linearity:

2. *curvature at both ends,*

An indication that an indication that an indication
longer or shorter tails than the theoretical
distribution.

S-shaped S-shaped S-shaped first above S-shaped
line, line, indicates heavier line, indicates heavier line, indi
distribution.

S-shaped S-shaped first below then above the
line, line, indicates line, indicates lighter tails than the t
distribution.

Displays of EDA

Departures from linearity:

3. *convex or concave curvatures,*

An indication the theoretical symmetric and the empirical one is not.

C-shaped (concave) below the $y=x$ line indicates positively skewed data.

C-shaped (convex) above the $y=x$ line indicates negatively skewed data.

Displays of EDA

Departures from linearity:

4. *horizontal segments, plateaus, or gaps*

Granularity in the data which occurs at certain values may be due to rounding (horizontal segments). Plateaus or gaps may be an indication of more than one theoretical distribution.

Displays of EDA

Cautions for interpreting theoretical Q - Q plots:

- 1.1. The natural variability of the distributional model from straightness.

Displays of EDA

Cautions for for interpreting theoretical for interpreting theor

2. Each *Q-Q plot* on only only compares the on
distribution of one distribution of one distr
distribution; distribution; distribution; all distribution; all d
data set, in particular the relationship data set, in
variable to others is ignored.



Conclusion: Theoretical *Q-Q pl**Q-Q plots* are not a
panacea and must
other displays and analyses to get a fu
the behavior of the data.

CDA Introduction

What is *confirmatory data analysis* (CDA)?

The role of CDA is closer to that of traditional statistical inference. It provides statements of significance and confidence, for example, inferential goodness-of-fit tests and tests for goodness-of-fit. Its function is to provide the statistician with insight into a set of data prior to evaluation and drawing conclusions.

CDA methods are used to assess

the *reproducibility* of observed patterns or effect of observed and

goodness-of-fit using statements of confidence and significance.

CDA Introduction

Goodness-of-fit hypothesis that a given stated probability law $F(x)$.

The null hypothesis can be a *simple hypothesis*

when $F(x)$ is completely specified, for example, normal with mean, $\mu = 10$; or

the null hypothesis can be when $F(x)$ is not completely specified, is not completely specified and .

CDA Introduction

Goodness-of-fit tests for nGoodness-of-fit tests for nGoodness-of-fit tests for n
into five categories:

chi-square tests,

empirical distribution function (EDF) tests,

moment tests,

regression tests, and

miscellaneous tests.

CDA Goodness-of-fit Tests

It is very hard to compare goodness-of-fit tests to establish criteria as to what is the best test for a particular situation. The particular situation that the alternative hypothesis for GOF is that the alternative hypothesis is vague, vague, vague, for example the empirical distribution is normal.

Comparison between GOF tests using *power* as the criteria.

Definition . . .

The *power* of a test, $1 - \beta$, is the probability of rejecting the null hypothesis when it is in the null hypothesis.

Make the DECISION:	The NULL HYPOTHESIS is:	
	True	False
<i>Not to Reject the Null Hypothesis</i>	<i>Correct Decision</i> ($1 - \alpha$)	<i>Incorrect Decision</i> Type II Error (β)
<i>to Reject the Null Hypothesis</i>	<i>Incorrect Decision</i> Type I Error (α)	<i>Correct Decision</i> Power ($1 - \beta$)

CDA Goodness-of-fit Tests

CChi-sChi-square type GOF tests were developed by Karl Pearson.

TheThe mechanics consistThe mechanics consist of hypothehypothethypothesizedhypothesized distribution (with parameters)parameters) into a multinomial distribution with cells,cells, cells, countingcells, counting cells, counting the cells, c each cell and contrasting these, using a chi-square or likelihoodlikelihood ratio test statistic, likelihood ratio number observations for each cell.

CDA *Goodness-of-fit* Tests

Some prominent chi-square G-tests are Fisher, Watson-Roy, and Rao-Robson.

Recommendations:

It is recommended that the chi-square GOF test not be used in testing departures from normality when the data are *complete*. (Complete data are those for which a value of each observation is observed. An incomplete or censored data would be an analysis measurement that was measured at a limit. Here we don't have complete data that the value is less than the value.) Other procedures to be discussed are more powerful (D'Agostino, 1986).

CDA *Goodness-of-fit* Tests

Empirical distribution function the discrepancy between the EDF and the distribution function, and are used to test the sample to the distribution. The sample to the distribution can be completely specified or may contain parameters which must be estimated from the sample.

The EDF is $F_n(y)$ defined by

For any x , $F_n(y)$ records the proportion of observations less than or equal to x . $F_n(y)$ is used to estimate $F(y)$. In fact it is a consistent estimator of $F(y)$, since as $n \rightarrow \infty$, $|F_n(y) - F(y)|$ decreases to zero with probability 1.

CDA Goodness-of-fit Tests

The EDF is just another distribution of a random variable. The EDF is an empirical *cumulative relative frequency* which is a simple example of a *cumulative distribution function* (CDF).

Definitions . . .

Frequency is the number of observations in a particular class. The *relative frequency* is expressed as a proportion. The *frequency distribution function* (PDF) is the relative frequency histogram.

The *cumulative frequency* is the number of observations less than or equal to a given observation. The *relative cumulative frequency* is the cumulative frequency expressed as a proportion or percent of the total frequency.

CDA Goodness-of-fit Tests

Examples of a simple PDF and CDF.

CDA Goodness-of-fit Tests

EEDFEDF tests are based on the largest vertical difference between $F_n(y)$ and $F(y)$. They are divided into three classes, *supremum* and *quadratic*.

Supremum:

The most well-known EDF test was introduced by Kolmogorov in 1933. It is referred to as the Kolmogorov-Smirnov test or the KS Test. D is the largest of two differences:

1. $F_n(y) > F(y), D^+ = \sup_y \{ F_n(y) - F(y) \},$ and
2. $F_n(y) < F(y), D^- = \sup_y \{ F(y) - F_n(y) \}.$

Combined we have,

$$D = \sup_y | F_n(y) - F(y) | = \max\{ D^+, D^- \}.$$

CDA Goodness-of-fit Tests

Graphical Representation of the KS Test

CDA Goodness-of-fit Tests

Quadratic:

The second class of EDF gives weights, the squared differences $[F_n(y) - F(y)]^2$.

One example is the Cramer-von Mises W^2 . For the Cramer-von Mises statistic, $F(y) = 1$.

Another example is the Anderson-Darling A^2 . For the Anderson-Darling statistic, $[(F(y))(1 - F(y))]^{-1}$.

CDA Goodness-of-fit Tests

Recommendations:

The most powerful EDF test is the Anderson-Darling, A^2 . Power studies are providing information on comparisons to non-EDF tests.

For testing normality, it is recommended to use the Kolmogorov-Smirnov (K-S) test only as a historical curiosity. The K-S test has less power in comparison to other procedures (D'Agostino, 1986).

CDA Goodness-of-fit Tests

Moment type GOF tests can be regarded as type GOF tests that have been initiated by Karl Pearson. He recognized that deviations from normality could be characterized by the standard third and fourth moments of the distribution.

CDA Goodness-of-fit Tests

A test of the third standardized moment $\sqrt{\frac{\mu_3}{\mu_2^3}}$,

$$H_0: \sqrt{\frac{\mu_3}{\mu_2^3}} = 0.$$

S_U Approximation (D Agostino, 1970)

(1) Compute $\sqrt{\frac{\mu_3}{\mu_2^3}}$ from the sample data.

(2) Compute

$$\sqrt{\frac{\left[\frac{\sum_{i=1}^n (x_i - \bar{x})^3}{n} \right]}{\left[\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n} \right]^3}}$$

$$\sqrt{\frac{\mu_3}{\mu_2^3}}$$

CDA Goodness-of-fit Tests

S_U Approximation (cont.)

(3) Compute

Z is approximately a standard normal variate.

This transformation is applicableThis transformation is applicable
greater than or equal to eight. test is excellent
for detecting nonnormality due to skewness.

CDA Goodness-of-fit Tests

A test of the fourth standardized moment b_2 ,

$$H_0: \mu_2 = 3.$$

Anscombe and Glynn Approximation (1983)

- (1) Compute b_2 from the sample data
- (2) Compute the mean and variance of b_2

- (4) Compute the third standardized moment of b_2

$$\sqrt{\frac{b_2 - \mu_2}{\sigma_2}} \quad \frac{b_2 - \mu_2}{\sigma_2} \quad \sqrt{\frac{b_2 - \mu_2}{\sigma_2}}$$

CDA Goodness-of-fit Tests

Anscombe and Glynn Approximation (cont.)

(5) Compute

$$\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n-1} \left[\frac{\sum_{i=1}^n (x_i - \bar{x})^3}{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\frac{n-1}{2}} \right]$$

(6) Compute

$$\left[\left(\frac{\sum_{i=1}^n (x_i - \bar{x})^3}{\sum_{i=1}^n (x_i - \bar{x})^2} \right) \sqrt{\frac{n-1}{2}} \right] \sqrt{\frac{n-1}{2}}$$

Z is approximately a standard normal variate.

The b_2 test is primarily used to detect nonnormal kurtosis. The b_1 test is primarily used to detect nonnormal skewness. The b_3 test is primarily used to detect nonnormal thickness.

A number of researchers have used these tests to produce an *omnibus* test of normality.

CDA Goodness-of-fit Tests

One *omnibus* test of normality is the *R*-test.

The *R*-test is the simplest omnibus test, it consists of performing the

χ^2 test at the α_1 level of significance, and the b_2 test at the α_2 level of significance.

The overall level of significance employed is the Bonferroni inequality, $\alpha_1 + \alpha_2$.

The term *R*-test was given to this procedure because it can be viewed as employing rectangular coordinates for the rejection of normality.

CDA Goodness-of-fit Tests

Another omnibus test of normality is the K^2 test of D'Agostino and Pearson (1973).

D'Agostino and Pearson suggested the test statistic,

$$\sqrt{\frac{K^2}{n}}$$

as an omnibus test where K^2 is a standardized normal variable with 2 degrees of freedom.

Recommendations:

The K^2 test is more powerful than the R -test.

CDA Goodness-of-fit Tests

Regression and *correlation* type type GOF tests make use of order order statistics, order statistics, $y_{(i)}$. A. A straight line is a *Q-Q plot* and GOF tests are constr and GOF tests are co statistics associated with the line,

$$E(y_{(i)}) = \quad + \quad m_i, \quad (1)$$

where where is a location parameter, is a scale parameter, parameter, a parameter, and m_i represents t distribution.

There There are three main appr There are three main approa the data fit equation (1).

1. A test based on the correlation coefficient.
2. A test based on the sum of squared residuals $\{ \quad \}, \}$, where \quad . In order to provide a scale-free test provide a scale-free test divided by another quadratic form.
3. The The scal The scale parameter The scale parameter, squared squared values squared value compared with another

2.
.

CDA Goodness-of-fit Tests

The Shapiro-Wilk GOF test is based on The Shapiro-Wilk method of testing the fit of model (1).

The steps for conducting the Shapiro-Wilk GOF test are provided below.

1. Calculate

, where

$r = (n - 1)/2$ if n is odd and $r = n/2$ if n is even, and a_i s are the optimal weights for least squares estimator of μ , given that t population is normally distributed.

2. Calculate

$$W = Y^2 / S^2 .$$

3. If W is less than the value in the lower tail table for the percentage points of the W -test for normality, for a particular null.

CDA Goodness-of-fit Tests

The exact distribution of W depends on n . Since this distribution is not . Since Shapiro and Wilk provided Monte Carlo points for use with the test for points for use with the 50.

CDA Goodness-of-fit Tests

Recommendations (D Agostino, 1986):

Graphical analyses should *always* accompany a formal test for normality.

The Shapiro-Wilks W test and the D Agost test and Pearson K^2 test appear to be the most available. The Shapiro-Wilks test is probably overall most powerful.

The K-S test should never be used. It has low power in comparison to other procedures.

When testing for normality with the chi-square test should have poor power in comparison to other procedures.

References

- D Agostino, D Agostino, R. B., and Stephens, M. A. (1986). *Goodness-of-Fit Tests for Normality*. Marcel Dekker.
- D Agostino, R. B. (1986). Tests for the Normality of a Sample. In *Techniques*, D Agostino, R. B., and Stephens, M. A., Eds., pages 367 - 419.
- D Agostino, R. B., Belanger, A., and D Agostino, Jr., R. (1986). Powerful and Informative Tests of Normality. *Journal of the American Statistical Association* **81**, 410-415.
- Freedman, D. and Diaconis, P. (1981). On the histogram as a density estimator. *Zeitschrift fur Wahrscheinlichkeitstheorie und verwandte Gebiete* **57**, 453-476.
- Hand, D. J. (1999). Chapter 1. Introduction. In *Intelligent Data Analysis*, Hand D. J., Eds., pages 1-14.
- Kruskal, W. H. (1960). Some remarks on wild observations. *Technometrics* **2**, 1-3.
- Scott, D. W. (1979). On optimal and data-based histograms. *Biometrika* **66**, 605-610.
- Scott, D. W. (1992). *Multivariate Density Estimation. Theory, Practice, and Visualization*. New York: John Wiley and Sons.
- Tietjen, G. L. (1986). The Analysis of Variance. In *Techniques*, D Agostino, R. B., and Stephens, M. A., Eds., pages 497 - 522.
- Wainer, H. (1984). How to Display Data Badly. *The American Statistician* **38**, 137-146.